

Generalized Assignment Problem: Truthful Mechanism Design without Money

Salman Fadaei*, Martin Bichler

*Department of Informatics, Technical University of Munich,
salman.fadaei@gmail.com, bichler@in.tum.de*

Abstract

In this paper, we study a mechanism design problem for a strategic variant of the generalized assignment problem (GAP) in a both payment-free and prior-free environment. In GAP, a set of items has to be optimally assigned to a set of bins without exceeding the capacity of any singular bin. In the strategic variant of the problem we study, bins are held by strategic agents, and each agent may hide its compatibility with some items in order to obtain items of higher values. The compatibility between an agent and an item encodes the willingness of the agent to receive the item. Our goal is to maximize total value (sum of agents' values, or social welfare) while certifying no agent can benefit from hiding its compatibility with items. The model has applications in auctions with budgeted bidders.

Keywords: Mechanism Design without Money, Generalized Assignment Problem, Truthfulness, Approximation

1. Introduction

Truthful mechanism design without money under general preferences is a classic topic in social choice theory. Truthfulness ensures that no agent can be better off by manipulating its true preferences. When searching for truthful mechanisms *without money*, one has to look at restricted domains of preferences. The reason for this, is the Gibbard-Satterthwaite theorem which states that any truthful social choice function which selects an outcome among three or more alternatives has to be trivially aligned with the

*Corresponding author

preference of a single agent (namely, the dictator) [1, 2]. Thus, exploring domains for which there exist truthful mechanisms is of central importance in the field of social choice theory.

As an example for restricted domains, consider agents with single-peaked preferences. In this domain returning the *median* of the peaks determines a truthful social choice [3]. Another example is the two-sided matching, in which a set of men has a strict preference ordering over a set of women, and vice versa. A matching is an assignment of men to women where each side is assigned to only one element of the other side. The *deferred acceptance algorithm* finds a stable matching which is truthful for the proposing side, but not necessarily truthful for the other side [4].

One way to circumvent the impossibility result is relaxing the social choice function. Procaccia and Tennenholtz introduced the technique of welfare approximation as a means to derive truthful approximation mechanisms without money [5]. This type of approximation is not meant to handle computational intractability, but a method to achieve truthfulness by relaxing the goal of optimizing social welfare (approximating social welfare), and thus circumventing the Gibbard-Satterthwaite impossibility theorem. The approach is to maximize welfare without considering incentives, and refer to this as optimal value. Then it is said that a truthful mechanism returns (at most) an α -approximation of the optimal if its value is always greater than or equal to $1/\alpha$ times the optimal value ($\alpha \geq 1$). Several works, subsequent to the work of Procaccia and Tennenholtz, employ this technique [6, 7, 8]. We apply this technique to a novel strategic setting in the following.

1.1. Model

Consider a strategic variant of the generalized assignment problem termed GAP-BS in an environment which is both prior-free and payment-free. In GAP-BS, there are m items J and n bins (knapsacks) I . Each bin i has a capacity C_i and associates a value v_{ij} and a size w_{ij} to any item j . A feasible assignment may allocate a subset of items S to bin i such that $\sum_{j \in S} w_{ij} \leq C_i$. A feasible assignment may assign each item at most once.

In GAP-BS, we assume tuple $T = (\{v_{ij}\}_{ij}, \{w_{ij}\}_{ij}, \{C_i\}_i)$ is public, but each bin is held by a strategic agent. The private information that each agent/bin holds is the set of its compatible items. The compatibility between an agent and an item encodes the willingness of the agent to receive the item. In particular, consider a bipartite graph G where one side corresponds to items and the other side corresponds to bins. The edges of G , $E \subseteq$

$I \times J$ represent the compatible item-bin pairs. The private type of a bin i is therefore the set of edges in the graph incident on i , i.e. E_i . A bin i receives value $v_i(S) = \sum_{j \in S: (i,j) \in E_i} v_{ij}$ from package S if $\sum_{j \in S} w_{ij} < C_i$ and 0, otherwise. The total value of a feasible assignment (S_1, S_2, \dots, S_n) equals the sum of values received by the bins from the assignment: $\sum_{i \in I} v_i(S_i)$. We seek a total value-maximizing algorithm that provides each bin i with incentives to truthfully report its compatible items E_i rather than any $E'_i \subset E_i$.¹ In other words, given a truthful mechanism, bins have no incentive to hide their compatibility with some items.

Let \mathcal{A} denote a randomized algorithm which takes instance (T, E) and computes $X \in \{0, 1\}^E$, an assignment of items to bins. Notice, the assignment itself is a deterministic assignment (each bin receives a deterministic set of items), but algorithm \mathcal{A} is internally randomized, i.e., \mathcal{A} returns a solution which is randomly chosen according to a probability distribution over feasible assignments. Thus, the computed assignment may change by running \mathcal{A} , twice on the same input. Randomized algorithm \mathcal{A} , given any tuple T , should satisfy the following properties.

- i. (feasibility) $\forall j \in J, \Pr[\sum_{i \in I} X_{ij} \leq 1] = 1$ and $\forall i \in I, \Pr[\sum_{j \in J} w_{ij} X_{ij} \leq C_i] = 1$, where $X \sim \mathcal{A}(T, E)$, for all E .
- ii. (incentive compatibility, or truthfulness) for all i , E_i , E_{-i} , and any reported $E'_i \subset E_i$, we have $\mathbb{E}[\sum_{j: (i,j) \in E_i} v_{ij} X_{ij}] \geq \mathbb{E}[\sum_{j: (i,j) \in E_i} v_{ij} X'_{ij}]$, where $X \sim \mathcal{A}(T, E)$, and $X' \sim \mathcal{A}(T, E'_i \cup E_{-i})$.

E_{-i} always denotes $E \setminus E_i$. The expectation in *ii* is taken over the coin flips of the algorithm. Note that, the expected value of the bin in both cases is calculated with respect to true item-bin compatibilities, E . We remark that condition *ii* characterizes mechanisms that are dominant strategy incentive compatible. In this paper, for brevity, we refer to these mechanisms as truthful mechanisms or algorithms. To sum, our objective is to propose a randomized algorithm \mathcal{A} for GAP-BS which is truthful, and always returns a feasible assignment whose value approximates the optimal total value as high as possible.

¹In fact, our results certify that each bin i reports exactly E_i , and has no incentives to report any other set of edges E'_i . However, for the sake of simplicity in the exposition of the results, we focus on untruthful reports that are made by hiding some edges, $E'_i \subset E_i$.

Many real-world decision problems can be modeled by variants of knapsack problems, therefore we believe that our model can be applied broadly. As an example, we refer to the *maximum budgeted allocations* (MBA) problem [9]. In MBA, a set of indivisible items has to be assigned to a set of bidders. Each bidder i reports her willingness to pay b_{ij} for item j by bidding for the item, while she has a budget constraint B_i . Each bidder i on receiving a package S of items, pays $\sum_{j \in S} b_{ij}$. Each bidder i has the rigid constraint B_i on her payment. The goal in MBA is to find a distribution of items among the bidders which maximizes the total revenue (the sum of the payments by the bidders while respecting their budget constraints). MBA arises in auctions with budgeted bidders and has several applications [9].

In MBA, bidders want to get as much as they can without spending more than their budget. For instance, advertisers wish to maximize the impressions, clicks, or sales generated by their advertising, subject to budget constraints. Similarly, bidders who have no direct utility for leftover money (e.g. because the money comes from a corporate budget) will buy as much as possible. This types of bidders are called *value maximizers*, and have recently drawn the attention of researchers in mechanism design [10, 11].

Consider a strategic variant of MBA in which each bidder, in order to obtain a more valuable package of items, strategizes in the following way. Each bidder may strategically hide her interests in buying some items by not bidding for those items. In this setting, the auctioneer wishes to certify that each bidder truthfully reveals her willingness to buy items. In other words, a truthful mechanism in this setting will encourage participation of the bidders in the auction. We model this setting by GAP-BS in which each bidder is represented by a bin, budgets B_i by capacities C_i , the bids b_{ij} by the values of bins for the items v_{ij} , and the payment by a bidder i for item j by the weight of the item on the bin, w_{ij} . Thus, in this setting of GAP-BS, we have $v_{ij} = w_{ij}$ for all i and j . For this problem, since the value density of each item is the same over all bins, we provide a truthful 4-approximation algorithm.

1.2. Discussion About the Assumptions

Aside from the applications of the model discussed above, we emphasize that our assumptions (which imply a highly structured domain) are necessary to escape the impossibility results such as the Gibbard-Satterthwaite theorem and its variations [12]. For example, we resort to welfare approximations because as stated by Theorem 1, no deterministic (or randomized) algorithm

whose value is optimal, exists for GAP-BS. The lower bounds in Theorem 1 were derived for a different setting with strategic items in the literature, however, we can reproduce and adapt the theorem for our setting.

Theorem 1. *[6] No truthful deterministic algorithm with an approximation ratio better than 2 exists for GAP-BS. Moreover, no truthful-in-expectation randomized algorithm with an approximation ratio better than 1.09 exists for GAP-BS.*

Now, we consider a setting in which bins/agents have private values for items. This setting is more general than GAP-BS in that, in this setting, the agents can manipulate their valuations for items. This is in contrast to GAP-BS in which the agents can only hide their valuations for some items by hiding their compatibility with those items. For this general setting, no deterministic (or randomized) truthful algorithm, with an interesting approximation ratio, exists. To see this, consider a simple market with one item, and a set of agents. This market is equivalent to the single-item auction, but without money. We observe that no mechanism without money can find the (true) highest valuation for the item, as the agents can report arbitrarily high values for the item. That is, no truthful algorithm can do any better than the algorithm which allocates the item to the bin which is uniformly chosen at random. Such an algorithm provides a trivial approximation ratio of $1/n$, n being the number of agents.

In a parallel setting, Dughmi et al. [6] and Chen et al. [7] studied GAP in an environment in which *items* are held by strategic agents. This is in contrast to our assumption that *bins* are held by strategic agents. Hence, the solutions proposed by these authors are not directly applicable to GAP-BS. In GAP each item can be assigned only once, thus the setting studied by Dughmi et al. is appropriate for modeling single-demand bidders who are interested in buying only a single item. However, our model analyzes strategic bins which can model multi-demand bidders, i.e., bidders who are interested in buying multiple items. In particular, the sizes in our model are at the side of strategic agents which properly models the bidders' budgets in MBA problem.

1.3. Results and Technique

In addition to GAP-BS, we also analyze two variants, namely the multiple knapsack problem in which each item has the same size and value over bins,

and density-invariant GAP in which each item has the same value density (value per size) over the bins.

We observe that the relaxation and rounding technique is applicable to these problems. The relaxation and rounding technique is a welfare approximation technique [5] based on linear programming relaxations. To apply the technique, we start with a linear programming relaxation of the problem. Then, we need an algorithm which returns a fractional solution to the relaxation with an acceptable approximation ratio. The algorithm has to be fractionally truthful, i.e., no agent can increase its fractional value by untruthful reports. Finally, a rounding scheme which preserves truthfulness is applied to the fractional solution to obtain an integer solution. It should be noted that the relaxation and rounding technique has been previously applied to mechanism design without money in a different setting [6]. This fact is realized in Theorem 2.

We apply the technique successfully to our problems by proposing fractionally truthful algorithms with acceptable approximation ratios. For the rounding scheme, we use a rounding method called *randomized meta-rounding*, originally proposed by Carr and Vempala [13], and later applied by Lavi and Swamy [14] to mechanism design (with quasi-linear valuations). Using the relaxation and rounding technique, for two variants of GAP-BS, the multiple knapsack problem, and density-invariant GAP, we propose truthful 4-approximation algorithms. For GAP-BS, we show an $O(\ln(U/L))$ -approximation mechanism where U and L are the upper and lower bounds for value densities of the compatible item-bin pairs.

2. Generalized Assignment Problem

We start with a linear programming relaxation of GAP-BS.

$$\begin{aligned}
& \text{Maximize} && \sum_{i=1}^n \sum_{j=1}^m v_{ij} x_{ij} && (\text{LP[E]}) \\
& \text{subject to} && \sum_{i=1}^n x_{ij} \leq 1 \quad \forall j \in J \\
& && \sum_{j=1}^m w_{ij} x_{ij} \leq C_i \quad \forall i \in I \\
& && x_{ij} \geq 0 \quad \forall i, j \\
& && x_{ij} = 0 \quad \forall (i, j) \notin E.
\end{aligned}$$

Our technique is as follows. We design a *fractionally truthful* approximation algorithm which returns a feasible solution to LP[E]. A fractionally

truthful algorithm allocates fractional assignments to bins, and no bin can improve its fractional value by an untruthful report. In particular, a fractionally truthful algorithm \mathcal{A}^F takes (T, E) and returns $x \in [0, 1]^E$, a feasible solution to $\text{LP}[E]$ with the following property. For each bin i , if the bin reports $E'_i \subset E_i$, we will have $\sum_{j:(i,j) \in E} v_{ij} x_{ij} \geq \sum_{j:(i,j) \in E} v_{ij} x'_{ij}$, where $x' = \mathcal{A}^F(T, E'_i \cup E_{-i})$ and $E_{-i} = E \setminus E_i$. Next, we round the fractional solution using a special rounding technique which makes sure that each bin obtains a fixed fraction of its fractional value in expectation. The *randomized meta-rounding* is capable of maintaining this fixed fraction.

To use the randomized meta-rounding, we have to scale down the fractional solution by factor 2, which is essentially the integrality gap of the $\text{LP}[E]$ [15]. Assuming $x^* = \mathcal{A}^F(T, E)$, the randomized meta-rounding represents $x^*/2$ as a convex combination of polynomially-many feasible integer solutions. Looking at the provided convex combination as a probability distribution over integer solutions, we sample a randomized solution X which is always feasible, and its expected value is $1/2$ of the fractional value of x^* . This is confirmed by Theorem 2 from the literature.

Theorem 2. [6] *If there exists a fractionally truthful α -approximation algorithm for GAP-BS, then there exists a truthful (2α) -approximation solution for GAP-BS.*

2.1. Multiple Knapsack Problem

We consider a variant of GAP-BS in which neither the size nor the value of each item depends on the bins. Formally, for each item j we have $v_{ij} = v_j$ and $w_{ij} = w_j$ for all bins i . First, we observe an algorithm that returns a (fractional) optimal solution to $\text{LP}[E]$ is not fractionally truthful. For more details, we provide Example 1 relegated to the Appendix.

We propose Algorithm 1. We choose bin i in an arbitrary order and (fractionally) assign compatible items to it according to the decreasing order of value densities of items v_j/w_j until the capacity of the bin is exhausted or all compatible items are exhausted. Then we proceed to the next bin with remaining (fractional) items.

Algorithm 1 is fractionally truthful. It is known that assigning items according to decreasing order of value densities, when fractional assignments are allowed, produces the highest fractional value for the bin. Since bins wish to maximize their values, and the algorithm is aligned with this goal, the bins have no incentive to lie, and the algorithm is fractionally truthful.

Algorithm 1: Multiple Knapsack Problem.

1. Sort items according to the decreasing order of value densities v_j/w_j , breaking ties arbitrarily.
2. **foreach** bin i *chosen in an arbitrary order* **do**
for each unassigned (fractional) item j where $(i, j) \in E$ in the order defined above, fractionally assign as much of the item to bin i until the item is exhausted or the bin is full.
end
- return** the resulting assignment x .

For example, consider a bin with capacity $M \gg 1$, and two items 1 and 2 such that $v_1 = 1 + \epsilon$, $w_1 = 1$, and $v_2 = w_2 = M$. Algorithm 1 assigns item 1, and $\frac{M-1}{M}$ -fraction of item 2 to the bin, resulting in $1 + \epsilon + \frac{M-1}{M}M = M + \epsilon$ value for the bin, the highest possible value. If fractional allocations were not allowed, only item 1 could be assigned to the bin since the remaining capacity ($M - 1$) is not sufficient to assign item 2. Thus, in the absence of fractional allocations, the bin has incentives to hide its compatibility with item 1 in order to obtain item 2.

With regard to the total value, we show that Algorithm 1 returns a 2-approximate fractional solution. We compare the outcome of the algorithm with the optimal solution to the LP formulation of the problem shown below.

$$\begin{array}{ll}
\text{Maximize} & \sum_{i=1}^n \sum_{j=1}^m v_j x_{ij} \\
\text{subject to} & \sum_{i=1}^n x_{ij} \leq 1, \quad \forall j \in J \\
& \sum_{j=1}^m w_j x_{ij} \leq C_i, \quad \forall i \in I \\
& x_{ij} \geq 0, \quad \forall i, j \\
& x_{ij} = 0, \quad \forall (i, j) \notin E.
\end{array} \tag{MKP-LP[E]}$$

Assignment x computed by Algorithm 1 is a feasible assignment since each item is assigned only once, and the capacity of the bins are respected by the algorithm. Thus, x belongs to the region of feasible solutions to MKP-LP[E].

Lemma 1. *Algorithm 1 returns a 2-approximation solution to MKP-LP[E].*

Proof. We will construct a feasible dual solution with a value at most twice the value obtained by the algorithm, then by calling the weak duality theorem, the claim will follow. Assume x is the outcome of Algorithm 1. Using x

we can construct a feasible solution to the dual of MKP-LP[E] given below.

$$\begin{array}{ll}
\text{Minimize} & \sum_{j=1}^m p_j + \sum_{i=1}^n u_i C_i & (\text{MKP-LPD}[E]) \\
\text{subject to} & p_j + u_i w_j \geq v_j, & \forall (i, j) \in E \\
& u_i \geq 0, & \forall i \\
& p_j \geq 0, & \forall j.
\end{array}$$

Initially, let $p = \vec{0}$ and $u = \vec{0}$. If item j gets exhausted, set $p_j = v_j$. Furthermore, for all *full* bins i , set $u_i = v_j/w_j$, j being the last item (fractionally) assigned to i . We can observe that this satisfies the constraint corresponding to each edge (i, j) . In particular, if bin i is full, then for each j incident on i , either j gets exhausted with this assignment or does not. If j is exhausted we have $p_j = v_j$ and therefore the constraint holds. If j is not exhausted, we have $v_j/w_j \leq u_i$ since items are assigned in decreasing order of value density and thus the constraint holds. If bin i is not full, every item j which is assigned to it is exhausted by this assignment. That is we have $p_j = v_j$ and the constraint thus holds. For every item j which is not assigned to the bin but $(i, j) \in E$, we have $p_j = v_j$ since the item is exhausted due to another assignment. In sum, we have constructed a feasible dual solution using x .

Now, we bound the value of the dual solution with respect to the primal solution. First, we observe that $\sum_{i,j} v_j x_{ij} \geq \sum_j p_j \sum_i x_{ij}$, since $p_j = v_j$ if j is fully exhausted and $p_j = 0$, otherwise. Second, $\sum_{i,j} v_j x_{ij} = \sum_i \sum_j \frac{v_j}{w_j} (w_j x_{ij}) \geq \sum_i u_i \sum_j (w_j x_{ij})$, since if $x_{ij} > 0$ then $v_j/w_j \geq u_i$. Therefore, we obtain

$$\begin{aligned}
2 \sum_{i,j} v_j x_{ij} &\geq \sum_j p_j \sum_i x_{ij} + \sum_i u_i \sum_j (w_j x_{ij}) \\
&= \sum_j p_j + \sum_i u_i C_i
\end{aligned}$$

Notice, only for items j which get exhausted ($\sum_i x_{ij} = 1$) we have $p_j > 0$ and only for full bins ($\sum_j w_j x_{ij} = C_i$) we have $u_i > 0$. The final term is the value of the dual, the desired conclusion. \square

Finally, we call Theorem 2 and obtain the following.

Theorem 3. *There exists a truthful 4-approximation mechanism for the multiple knapsack problem in our model.*

2.2. Truthful Mechanism for GAP-BS

Now, we attempt to design a truthful algorithm for GAP-BS, but first solve the problem with an additional assumption. We assume that the *value density* of each item is the same over all bins. More formally, there exists a value d_j for each item j such that for all bins i , we have $\frac{v_{ij}}{w_{ij}} = d_j$. This assumption will be relaxed in Subsection 2.3. We design a truthful 4-approximation mechanism for GAP-BS under this extra assumption.

The proposed algorithm can be viewed as a variant of the *deferred acceptance algorithm* designed for matching marketplaces. Each item j has a preference list \mathcal{L}_j according to decreasing order of v_{ij} where $(i, j) \in E$, breaking ties arbitrarily. The preference list of a bin is defined according to the decreasing order of value densities. Once a (fractional) item and a bin are matched, the assignment will never be broken.

Algorithm 2: GAP with Equal Density
<p>Data: Preference lists of the items, $\{\mathcal{L}_j\}_j$.</p> <p>Result: A feasible solution x to $\text{LP}[E]$.</p> <ol style="list-style-type: none"> 1. Sort items according to their decreasing order of value densities d_j, breaking ties arbitrarily. 2. foreach item j <i>chosen according to the order above</i> do Fractionally assign as much of the item to the bins chosen according to the order specified by \mathcal{L}_j, until the item is exhausted or all the bins in \mathcal{L}_j are full. <p>return the resulting assignment x.</p>

To show the approximation factor of the solution, we can construct a feasible dual solution whose value is at most twice the value obtained by Algorithm 2, then by calling the weak duality theorem, the following lemma holds.

Lemma 2. *Algorithm 2 returns a 2-approximation solution to $\text{LP}[E]$ when each item has the same value density over bins.*

The truthfulness proof of Algorithm 2 proceeds as follows. We first show that in an instance with 2 items and 2 bins (2×2), truthfulness holds. This instance contains the core of the truthfulness proof for the general case. Truthfulness for simpler cases is trivial. A straightforward generalization of the argument for 2×2 shows truthfulness for settings with m items and 2

bins $(2 \times m)$ for any $m > 2$. For the general case of $(n \times m)$ we provide an inductive argument.

In order to show that Algorithm 2 is *fractionally truthful*, we look at Algorithm 2 as a variant of the deferred acceptance algorithm where items propose capacities to bins. In Step 2 of the algorithm, we process items one by one. For each item, we try to assign the item or part of the item to the bins according to the decreasing order of the item values for the bins. To simplify the exposition of the proof, we view this process as items proposing to the bins. When bin i reveals its compatibility with item j , we view it as bin i accepting (possible) proposal by item j as far as the capacity of the bin permits. Similarly, bin i hiding its compatibility with item j can be viewed as bin i rejecting (possible) proposals by item j , or equivalently not allowing item j to propose to bin i .

Lemma 3. *Algorithm 2 is fractionally truthful for 2×2 settings.*


Proof. Let $\underline{1}$, and $\underline{2}$ denote the bins, and \underline{p} and \underline{q} denote the items. Let us assume \underline{p} precedes \underline{q} in proposing to the bins, i.e. $d_p \geq d_q$. Fix this order of proposing items as well as the reports by bin $\underline{2}$. We argue that bin $\underline{1}$ is never better off by hiding some of its edges E_1 . Showing the truthfulness for bin $\underline{2}$ is analogue.

Assume $(\underline{1}, \underline{q}) \in E_1$. Then bin $\underline{1}$ may receive a proposal from \underline{q} , but obviously the bin receives no proposal from \underline{q} if the bin reports $(\underline{1}, \underline{q}) \notin E_1$. Thus, hiding compatibility with \underline{q} might only make a loss for the bin.


Now, we analyze the behavior of the algorithm for a similar change in report for item \underline{p} . We need to show that when $(\underline{1}, \underline{p}) \in E_1$ (case I) the obtained value by the bin is at least as good as when $(\underline{1}, \underline{p}) \notin E_1$ (case II). Then we conclude that when truly $(\underline{1}, \underline{p}) \in E_1$, the bin has no incentive to report $(\underline{1}, \underline{p}) \notin E_1$.

In case I, if only a fraction of the proposal by \underline{p} is accepted by the bin, then the bin has become full by accepting a fraction of \underline{p} (recall \underline{p} precedes \underline{q} in proposing to the bins). Thus the obtained value by the bin is maximum and can't be better off in case II. If in case I no fraction of the proposal by \underline{p} is accepted by bin $\underline{1}$, or if there is no proposal by \underline{p} then there will be no improvement in the value of the bin in case II, as well. What remains is to show that the bin cannot be better off in case II when it accepts the proposal by \underline{p} fully in case I.

This situation is depicted in Figure 2.1. In the figure, (a) and (b) correspond to case I and case II, respectively. In the figure, C_{ij} denotes the capac-



(a) Case I. \underline{p} is exhausted when it is assigned to $\underline{1}$.



(b) Case II. $\underline{1}$ hides its compatibility with \underline{p} . At least a fraction of \underline{q} is assigned to $\underline{1}$ ($C'_{1q} > 0$).

Figure 2.1: Two cases where the bin is and is not on the preference list of the item. The amount of proposed and accepted capacities are shown on the edges.

ity proposed by item j to bin i , which is accepted by the bin. Considering the information provided in Fig. 2.1, we need to show that $C'_{1q} \leq C_{1q} + C_{1p}$. This will mean, in case II, the bin actually receives less capacity from items with less (or equal) value densities than in case I, which in turn means a lower value for bin $\underline{1}$. Notice, to arrive at this inequality we used the assumption that the order of proposing items is fixed in the two setups. To show the inequality, we first observe two facts about Algorithm 2.

Observation 1. *If a set of items together propose a capacity of $C^0 \leq C$ to a bin with capacity C , the bin will accept the whole proposed capacity. If we first let a capacity C^1 propose to the bin and afterwards let the foregoing items propose the capacity C^0 , the bin will reject a capacity of at most C^1 from the items that propose after the first capacity.*

Proof. Assume C^1 and C^0 in order propose to the bin. If the bin gets full by accepting C^1 we must have $C^1 \geq C$, then the bin will reject exactly a capacity of C^0 of the next items. Now because $C^0 \leq C \leq C^1$, the claim holds. If not (the bin still has an empty capacity of C^E after accepting C^1),

the bin accepts C^1 fully and rejects an amount equal to $\max\{0, C^0 - C^E\}$ from the next proposing capacities. We have $C^1 + C^E = C \geq C^0$, therefore $C^0 - C^E \leq C^1$. Thus, in this case the rejected capacity will be upper bounded by C^1 . This completes the proof. \square

Observation 2. *Let $\underline{1}$ and $\underline{2}$ be two subsequent bins in \mathcal{L}_j . If bin $\underline{1}$ rejects the proposed capacity C_{1j} by item j then, this is an upper bound to C_{2j} , the capacity that will be proposed by item j to $\underline{2}$, i.e. $C_{1j} \geq C_{2j}$.*

Proof. First, we must have $w_{1j} \geq w_{2j}$ since $\frac{v_{1j}}{w_{1j}} = \frac{v_{2j}}{w_{2j}}$ by the assumption of equal density over bins and $v_{1j} \geq v_{2j}$ as $\underline{1}$ precedes $\underline{2}$ in \mathcal{L}_j . Rejecting C_{1j} means that this fraction of the item remains: C_{1j}/w_{1j} . Then what will be proposed to $\underline{2}$ is $C_{2j} = w_{2j} \cdot (C_{1j}/w_{1j}) \leq C_{1j}$. This completes the proof. \square


Back to the argument about cases I and II, we notice that in case II there is an increase of amount C_{2p} in the proposed capacity to $\underline{2}$ compared to case I. The capacity rejected by bin $\underline{2}$ is thus upper bounded by C_{2p} according to Observation 1. That means, $C_{2q} - C'_{2q} \leq C_{2p}$. Moreover, according to Observation 2, the rejected capacity upper bounds the proposed capacity to the next bin. Hence, we have $C'_{1q} - C_{1q} \leq C_{2q} - C'_{2q}$. Therefore, we obtain $C'_{1q} \leq C_{1q} + C_{2p} \leq C_{1q} + C_{1p}$. The last inequality holds again because of Observation 2 (see it as bin $\underline{1}$ rejecting C_{1p} , an upper bound to C_{2p}). This completes the proof of Lemma 3. \square

A simple generalization of the argument for 2×2 markets shows truthfulness for the $2 \times m$ markets with $m > 2$. A useful observation here is that we only need to show that bin $\underline{1}$ will always report E_1 rather than $E_1 \setminus \{e_j\}$ for every $e_j \in E_1$. If we show this, we have in fact shown that reporting E_1 is better than reporting $E_1 \setminus \{e_j\}$. This also shows that reporting $E_1 \setminus \{e_j\}$ is better than hiding one edge from $E_1 \setminus \{e_j\}$, i.e. reporting $E_1 \setminus \{e_j, e_{j'}\}$ and so on. For the general case we provide an inductive argument. We assume that in a $(n - 1) \times m$ setting bins are truthful and prove that in a $n \times m$ setting truthfulness holds as well.


Lemma 4. *If Algorithm 2 is truthful for markets with m items and $n - 1$ bins, it will be truthful for $n \times m$ markets for $n \geq 3$, and $m \geq 2$.*

Proof. Consider bin \underline{i} and fix the reports of other bins denoted by $\underline{-i}$. We assume $(\underline{i}, \underline{p}) \in E_i$ (case I) and show that the bin will never be better off by reporting $(\underline{i}, \underline{p}) \notin E_i$ (case II). We compare the utility of the bin in the two

cases under a fixed order of proposing items. The two cases are depicted in Figure 2.2. Since the items before \underline{p} are assigned similarly in both cases, we only consider the items which are processed after \underline{p} denoted by $\underline{-p}$.



(a) Case I. \underline{p} is exhausted when it is assigned to \underline{i} . \underline{i} may get a fraction or nothing from other items $\underline{-p}$.



(b) Case II. \underline{i} hides its compatibility with \underline{p} . At least a fraction of $\underline{-p}$ is assigned to \underline{i} . \underline{p} is (fully) accepted by $\underline{-i}$.

Figure 2.2: Two cases where the bin shows or hides its compatibility with an item.

We show that $C'_{i,-p} \leq C_{ip} + C_{i,-p}$, where $C_{i,-p} = \sum_{q \in -p} C_{i,q}$ and $C'_{i,-p} = \sum_{q \in -p} C'_{i,q}$. This means that bin \underline{i} in case II actually receives less capacity from items with less (or equal) value densities than case I, which in turn implies lower value for the bin.

Consider case II. We look closer at the bin(s) to which item \underline{p} will be assigned. We assume \underline{p} is (fractionally) assigned to at least one bin otherwise we have $C_{i,-p} = C'_{i,-p}$ and thus the claim holds. Let bin $\underline{1}$ be the first bin to which \underline{p} will be assigned.

We assume bin $\underline{1}$ gets full at some point otherwise this bin accepts the extra capacity (C_{1p} , the capacity proposed by item \underline{p} to bin $\underline{1}$) without rejecting any capacity and therefore we have $C_{i,-p} = C'_{i,-p}$ and thus the claim holds. When bin $\underline{1}$ gets full, some of the currently proposing items to bin $\underline{1}$ will stop proposing to it and go to the next bin in their preference list. Let us call these capacities C_1 . C_1 is upper bounded by C_{1p} according to

Observation 1 which in turn is upper bounded by C_{ip} based on Observation 2: $C_1 \leq C_{1p} \leq C_{ip}$. If C_1 directly proposes to bin \underline{i} , the bin won't be better off in case II because $C_1 \leq C_{ip}$. The situation is worse for bin \underline{i} , if C_1 goes to the other bins. One can view this situation as bin \underline{i} rejecting capacity C_1 in a $(n-1) \times m$ setting where bin $\underline{1}$ (which is now full) and its absorbed capacities are eliminated. According to our induction assumption, this strategy will not make bin \underline{i} better off in a $(n-1) \times m$ setting. This completes the proof. \square

Taking into account, Lemma 3 and Lemma 4, we obtain the following.

Lemma 5. *Algorithm 2 is fractionally truthful.*

Finally, by calling Theorem 2, we obtain the following.

Theorem 4. *There exists a truthful 4-approximation mechanism for GAP-BS when each item has the same value density over all bins.*

2.3. Unequal Value Densities

We presented a truthful 4-approximate mechanism for GAP-BS when each item has a unique value density over all bins. Now we explain how to relax this assumption at the expense of a logarithmic loss in the total value. Consider those edges in E , $e = (k, l)$ and $e' = (k', l')$ whose value densities are respectively upper and lower bounds over all value densities:

$$L = \frac{v_{k'l'}}{w_{k'l'}} \leq \frac{v_{ij}}{w_{ij}} \leq \frac{v_{kl}}{w_{kl}} = U, \quad \forall (i, j) \in E.$$

Let us assume U and L are publicly known. This assumption will be removed later. Knowing this information we choose a density value d uniformly at random from the set $D = \{U, \frac{U}{2}, \frac{U}{4}, \dots, \frac{U}{2^{\mathcal{O}(\ln(U/L))}}\}$. Then we define a new valuation \hat{v} as follows. For every edge (i, j) in E with $\frac{v_{ij}}{w_{ij}} < d$ we set $\hat{v}_{ij} = 0$, or equivalently the edge is discarded from the graph. For every $\frac{v_{ij}}{w_{ij}} \geq d$, define \hat{v}_{ij} such that $\frac{\hat{v}_{ij}}{w_{ij}} = d$. Notice that always $\hat{v}_{ij} \leq v_{ij}$. Now we have an instance of GAP-BS with equal densities for which there exists a truthful 4-approximate mechanism according to Theorem 4. To ensure truthfulness, in the end, if item j is assigned to bin i by the subroutine for equal value densities, we withdraw the item with probability $1 - \frac{\hat{v}_{ij}}{v_{ij}}$. In other words, we let the bin hold the item with probability $\frac{\hat{v}_{ij}}{v_{ij}}$. If item j is assigned to bin i , the generated value for the bin will be v_{ij} , but if we let the bin hold the item

with probability $\frac{\hat{v}_{ij}}{v_{ij}}$, then the expected value will be \hat{v}_{ij} . This way, we make sure that each item has the same value density over all bins as it is required by the subroutine to guarantee truthfulness.

Set D contains $O(\ln(U/L))$ densities, and each density has the probability of $p = \frac{1}{O(\ln(U/L))}$ to be chosen. At least half of each valuation v_{ij} with probability p is counted in the expected total value; therefore, we obtain an $O(\ln(U/L))$ approximation factor.

To remove the assumption that U and L are public, we certify that the bins k and k' wouldn't hide the corresponding edges. To this end, we run one of the following three algorithms with probability $1/3$. *i*) Let bin k (the owner of edge e) choose all its desired items and assign nothing to the other bins. *ii*) Let bin k' (the owner of e') choose all its desired items and assign nothing to the other bins. *iii*) Exclude bin k and k' and run the algorithm above for all other bins using U and L obtained from the two excluded bins. One can observe that the two bins k and k' cannot do any better by hiding their edges. Also, it is easy to observe that the approximation factor is still $O(\ln(U/L))$. Thus, we obtain the following.

Theorem 5. *There exists a truthful $O(\ln(U/L))$ approximate mechanism for GAP-BS.*

We leave open the question of whether there exists a truthful mechanism with a constant factor of approximation for GAP-BS.

Acknowledgement

The authors gratefully acknowledge support of the Deutsche Forschungsgemeinschaft (DFG) (BI 1057/7-1). The authors also would like to acknowledge the anonymous referees for their useful comments and suggestions.

References

- [1] A. Gibbard, Manipulation of voting schemes: a general result, *Econometrica* 41 (1973) 587–601.
- [2] M. A. Satterthwaite, Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions, *Journal of economic theory* 10 (2) (1975) 187–217.
- [3] H. Moulin, On strategy-proofness and single peakedness, *Public Choice* 35 (4) (1980) 437–455.

- [4] A. E. Roth, M. A. O. Sotomayor, Two-sided matching: A study in game-theoretic modeling and analysis, no. 18, Cambridge University Press, 1992.
- [5] A. D. Procaccia, M. Tennenholtz, Approximate mechanism design without money, *ACM Transactions on Economics and Computation* 1 (4) (2013) 18.
- [6] S. Dughmi, A. Ghosh, Truthful assignment without money, in: *Proceedings of the 11th ACM conference on Electronic commerce*, ACM, 2010, pp. 325–334.
- [7] N. Chen, N. Gravin, P. Lu, Truthful generalized assignments via stable matching, *Mathematics of Operations Research* 39 (3) (2013) 722–736.
- [8] E. Koutsoupias, Scheduling without payments, *Theory of Computing Systems* 54 (3) (2014) 375–387.
- [9] D. Chakrabarty, G. Goel, On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP, *SIAM Journal on Computing* 39 (6) (2010) 2189–2211.
- [10] C. A. Wilkens, R. Cavallo, R. Niazadeh, Mechanism design for value maximizers, Cornell Working paper.
- [11] S. Fadaei, M. Bichler, Truthfulness and approximation with value-maximizing bidders, in: *the proceedings of the 9th International Symposium on Algorithmic Game Theory*, 2016.
- [12] S. Barbera, B. Peleg, Strategy-proof voting schemes with continuous preferences, *Social choice and welfare* 7 (1) (1990) 31–38.
- [13] R. Carr, S. Vempala, Randomized metarounding, in: *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, ACM, 2000, pp. 58–62.
- [14] R. Lavi, C. Swamy, Truthful and near-optimal mechanism design via linear programming, *Journal of the ACM (JACM)* 58 (6) (2011) 25.
- [15] D. B. Shmoys, É. Tardos, An approximation algorithm for the generalized assignment problem, *Mathematical Programming* 62 (1-3) (1993) 461–474.

Appendix A.

Proof of Theorem 1 *No truthful deterministic algorithm with an approximation ratio better than 2 exists for GAP-BS. Moreover, no truthful-in-expectation randomized algorithm with an approximation ratio better than 1.09 exists for GAP-BS.*

Proof. Consider a small market with two bins and two items shown in Figure A.1 (a). In this market, bins have capacity 1, and are both compatible with the two items. Item B is more valuable to both bins ($x > 1$), but each item has size 1. This market can be viewed as an instance for both the multiple-knapsack problem (Subsection 2.1), and the density-invariant GAP (Subsection 2.2).

Regarding deterministic mechanisms, an arbitrary truthful mechanism has to assign item B to one bin. Without loss of generality, assume B is assigned to bin $\underline{2}$, i.e., the tie is broken deterministically (alphabetically) in favor of bin $\underline{2}$. Now, consider reports in (b) of Figure A.1. The mechanism, in case (b), cannot assign B to bin $\underline{1}$ as it violates truthfulness. Thus, the mechanism assigns B to $\underline{2}$, and this results in an approximation ratio of $\frac{x+1}{x}$ which tends to 2 when x gets very close to 1.

An arbitrary truthful-in-expectation mechanism, in case (a), assigns B to one bin with a probability less than or equal $1/2$. Without loss of generality, let bin $\underline{1}$ be that bin. The utility of bin $\underline{1}$, in this case, will be at most $\frac{x+1}{2}$ for $x > 1$. Assume, in case (b), item B is assigned to bin $\underline{1}$ with probability q , resulting a $q \cdot x$ expected value for bin $\underline{1}$. Truthfulness stipulates no increase in the utility of bin $\underline{1}$ in case (b), i.e., $\frac{x+1}{2} \geq q \cdot x$, thus $q \leq \frac{x+1}{2x}$. In case (b), the total expected value will be $q(1+x) - (1-q)x = x+q$, thus the approximation ratio will be $\frac{x+1}{x+q}$. In order to obtain a smaller approximation ratio, we plug in $q = \frac{x+1}{2x}$. The ratio gets a value of $1 + \frac{1}{4\sqrt{2}+5} \approx 1.094$ for $x = 1 + \sqrt{2}$, the desired conclusion. \square

Proof of Theorem 2. *If there exists a fractionally truthful α -approximation algorithm for GAP-BS, then there exists a truthful (2α) -approximation solution for GAP-BS.*

Proof. Let \mathcal{A}^F denote a fractionally truthful algorithm for GAP-BS that takes an instance (T, E) and returns a feasible solution to $\text{LP}[E]$. Let x^* be the outcome of \mathcal{A}^F on instance (T, E) . Let $\{X^l\}_{l \in L}$ denote the set of feasible integer




Figure A.1: Circles represent items and squares represent bins. The $\frac{\text{value}}{\text{size}}$ of each item is on its left. Each bin has a capacity of 1. Selected assignments are in bold.

solutions to $\text{LP}[\mathbf{E}]$, where L indexes all feasible integer solutions. The integrality gap of $\text{LP}[\mathbf{E}]$ equals 2 [15], thus we scale down the fractional solution by factor 2. The meta-randomized rounding applied to $x^*/2$ computes a probability distribution over feasible integer solutions whose support is polynomial [13, 14]:

$$\frac{x^*}{2} = \sum_{l \in L} \lambda_l X^l, \quad \sum_{l \in L} \lambda_l = 1, \quad \text{and} \quad \forall l \in L, \lambda_l \geq 0.$$

We treat the convex decomposition above as a probability distribution according to which solution X^l has probability λ_l of being selected. Let X be a solution sampled from the above distribution. Obviously X is feasible by the construction of the distribution. We also have $\mathbb{E}[X_{ij}] = \frac{1}{2}x_{ij}^*$ for all i and j from the construction of the distribution. By the linearity of expectation, the expected value of a bin is $\mathbb{E}[\sum_{j:(i,j) \in E} v_{ij} X_{ij}] = \frac{1}{2} \sum_{j:(i,j) \in E} v_{ij} x_{ij}^*$. Therefore, the expected value of the solution is exactly half of the value of the fractional solution.

For truthfulness, fix bin i and $E_{-i} = E \setminus E_i$. Suppose the bin reports $E'_i \subset E_i$ rather than E_i . Let $x' = \mathcal{A}^F(T, E'_i \cup E_{-i})$, and X' be the solution

returned by the meta-randomized rounding from $x'/2$. We have

$$\begin{aligned}\mathbb{E}[\sum_{j:(i,j) \in E} v_{ij} X_{ij}] &= \frac{1}{2} \sum_{j:(i,j) \in E} v_{ij} x_{ij}^* \\ &\geq \frac{1}{2} \sum_{j:(i,j) \in E} v_{ij} x'_{ij} \\ &= \mathbb{E}[\sum_{j:(i,j) \in E} v_{ij} X'_{ij}]\end{aligned}$$

The inequality is because \mathcal{A}^F is fractionally truthful. Therefore, the bin cannot improve its expected value by hiding some of its edges. This completes the proof. \square

Proof of Lemma 2. Algorithm 2 returns a 2-approximation solution to LP[E] when each item has the same value density over bins.

Proof. An argument similar to that of Lemma 1 in addition to some required modifications will show the claim. Assume x is the outcome of Algorithm 2. Using x we can construct a feasible solution to the dual of LP[E] (LPD[E] given below) which is not greater than twice the value of x . Then we call the weak LP-duality theorem and conclude that x is a 2 approximate solution to LP[E].

$$\begin{aligned}\text{LPD[E]:} \\ \text{Minimize} \quad & \sum_{j=1}^m p_j + \sum_{i=1}^n u_i C_i \\ \text{subject to} \quad & p_j + u_i w_{ij} \geq v_{ij}, \quad \forall (i, j) \in E \\ & u_i \geq 0, \quad \forall i \\ & p_j \geq 0, \quad \forall j.\end{aligned}$$

Initially, let $p = \vec{0}$ and $u = \vec{0}$. If item j gets exhausted when assigned to bin i , set $p_j = v_{ij}$. Furthermore, for all *full* bins i , set $u_i = d_j$, j being the last item (fractionally) assigned to i . We can observe that this satisfies the constraint corresponding to each edge (i, j) . In particular, if bin i is full, then for each j incident on i , j either gets exhausted with this assignment or does not. If j is exhausted, we have $p_j = v_{ij}$ and therefore the constraint holds. If j is not exhausted, we have $v_{ij}/w_{ij} = d_j \leq u_i$ since items are assigned in decreasing order of value density and thus the constraint holds. If bin i is not full, every item j which is assigned to it is exhausted by this assignment. That is we have $p_j = v_{ij}$ and the constraint thus holds. For every item j which is not assigned to the bin but $(i, j) \in E$, we have $p_j \geq v_{ij}$ since the


item is exhausted due to an assignment $(i', j) \in E$ with $v_{i'j} \geq v_{ij}$. Therefore, we have constructed a feasible dual solution using x .

Now, we bound the value of the dual solution with respect to the primal solution. First, we observe that $\sum_{i,j} v_{ij}x_{ij} \geq \sum_j p_j \sum_i x_{ij}$, since p_j lower bounds the value of any edge on which any part of item j is assigned ($x_{ij} > 0$) because the item goes to bins according to the order specified by \mathcal{L}_j . Second, $\sum_{i,j} v_{ij}x_{ij} = \sum_i \sum_j \frac{v_{ij}}{w_{ij}}(w_{ij}x_{ij}) \geq \sum_i u_i \sum_j (w_{ij}x_{ij})$, since if $x_{ij} > 0$ then $\frac{v_{ij}}{w_{ij}} = d_j \geq u_i$. Therefore, we obtain


$$\begin{aligned} 2 \sum_{i,j} v_{ij}x_{ij} &\geq \sum_j p_j \sum_i x_{ij} + \sum_i u_i \sum_j (w_{ij}x_{ij}) \\ &= \sum_j p_j + \sum_i u_i C_i \end{aligned}$$

Notice, only for item j which gets exhausted ($\sum_i x_{ij} = 1$), we have $p_j > 0$ and only for full bins ($\sum_j w_j x_{ij} = C_i$) we have $u_i > 0$. The final term is the value of the dual, the desired conclusion. \square

Example 1 (Multiple Knapsack Example). *we observe an algorithm that returns a (fractional) optimal solution to $LP[E]$ is not fractionally truthful for the multiple knapsack problem. This can be seen in the example shown in Figure A.2. In (a) of this figure, the edges are reported truthfully, and the value-maximizing allocation, assigns A to bin $\underline{1}$ and the other item to the other bin. In (b) of this figure, bin $\underline{1}$ hides its compatibility with item A and as a consequence it is better off (in expectation) when the mechanism maximizes the total value. In (b) of this figure, the tie can be broken randomly or deterministically (alphabetically) in favor of bin $\underline{1}$. In any case, bin $\underline{1}$ is better off by manipulation.*



(a) True edges.



(b) Manipulated edges.

Figure A.2: Circles represent items and squares represent bins. The value/size of each item is on its left. Value maximizing assignments are in bold. $x \gg 1$.